arXiv:1411.4296v1 [cs.CV] 16 Nov 2014

Combining contextual and local edges for line segment extraction in cluttered images

Rui F. C. Guerreiro*

Institute for Systems and Robotics, Instituto Superior Técnico, Lisboa, Portugal

Abstract

Automatic extraction methods typically assume that line segments are pronounced, thin, few and far between, do not cross each other, and are noise and clutter-free. Since these assumptions often fail in realistic scenarios, many line segments are not detected or are fragmented. In more severe cases, *i.e.*, many who use the Hough Transform, extraction can fail entirely. In this paper, we propose a method that tackles these issues. Its key aspect is the combination of thresholded image derivatives obtained with filters of large and small footprints, which we denote as *contextual* and *local edges*, respectively. Contextual edges are robust to noise and we use them to select *valid* local edges, *i.e.*, local edges that are of the same type as contextual ones: dark–to–bright transition of vice-versa. If the distance between valid local edges does not exceed a maximum distance threshold, we enforce connectivity by marking them and the pixels in between as edge points. This originates connected edge maps that are robust and well localized. We use a powerful two-sample statistical test to compute contextual edges, which we introduce briefly, as they are unfamiliar to the image processing community. Finally, we present experiments that illustrate, with synthetic and real images, how our method is efficient in extracting complete segments of all lengths and widths in several situations where current methods fail.

Keywords:
Statistical two-sample tests, line segment detection, connectivity, connected segments, line pattern analysis, edge analysis

*Corresponding author
Email address: ruifcguerreiro@hotmail.com (Rui F. C. Guerreiro)

1. Introduction

Line segments provide important information about the geometric content of real-life images. Since most man-made objects are made of flat surfaces, the contours needed for the interpretation of such 2D images as well as 3D world scenes often consist of line segments. Also, many complex shapes accept an economic and simple description in terms of straight lines. This has been used, *e.g.*, for localizing vanishing points [1] or to match line segments across distinct views [2]. Other applications include, *e.g.*, rectangle detection [3], the inference of shape from lines [4], map-to-image registration [5], 3D reconstruction [6], or even image compression [7].

Although automatic line segment extraction has been researched actively in the past decades, current solutions make use of implicit strong assumptions that limit their applicability to simple and mostly unrealistic scenarios. Typical assumptions are that, *e.g.*, line segments are pronounced, thin, occur in small amounts, do not cross each other, are located away from noise or clutter. Some methods also assume that images have few data apart from line segments, such as textures or contours. Since these assumptions often fail in realistic scenarios, many line segments are not detected, or are fragmented to various extents. In more severe cases, extraction can fail altogether. For this reason, which we detail in the sequel, the robust detection of line segments in realistic scenarios remains an open frontier (see [8, 9, 10, 11] for examples of recent advances).

1.1. Overview of methods for line segment extraction

The Hough transform (HT) [12, 13] is the most popular method to detect lines in images. It is a likelihood-based parameter extraction technique that, basically, indicates that the largest accumulation of edge points correspond to image lines. It uses a Hough space, a two-dimensional space where each point represents a line in the image, and each edge point in the image votes on the region of the Hough space that represents the pencil of all the image lines that go through that edge point. By processing all edge points, the votes for each location in the Hough space are accumulated and the locations with larger number of votes correspond to the most likely parameterizations of the lines in the image. Later, the HT was extended for extracting line segments. After obtaining the line parameterizations with the usual HT, the start and end points of the segments are obtained using a gap-and-length

method [14], the shape of the spread of votes in the Hough space [15, 11] or extra accumulators [16].

The success of the HT comes from its global nature, since all points in a line contribute to its detection — in fact, it was proven that it implements a statistically robust estimator for finding lines [17]. It has, however, three major issues when used to extract line segments in complex images. Firstly, it requires an edge detection scheme such as, *e.g.*, the Canny edge detector [18], to generate its input edge map. Edge detection is by itself a hard problem, recognized as ill-posed in general [19], where delicate balances occur between edge localization and noise reduction, and detecting spurious edge points in noisy or textured areas and missing faint edges. As a consequence, edge detectors typically make use of small (local) filters and high thresholds for accepting an edge point, resulting in partial or complete segment mis-detections. This counteracts the global nature of the HT.

Secondly, since all votes originated by an edge point are *wrong* except for the one corresponding to the actual segment, the amount of noise in the Hough space is significant for images with many edge points [20]. This makes it difficult to identify the most likely parameterization of actual lines in the Hough space. This is particularly critical for short segments, since they originate small peaks that are hard to identify [21]. Accidental alignments of unrelated edge points can also originate false peak detections [20, 22]. An example where Hough space contamination due to poor edge detection and erroneous vote accumulations leads to complete extraction breakdown is shown in [22]. Although Duda and Hart [13] argued as early as 1972 that the noise in the Hough space can be reduced by taking connectivity between collinear points into account, this topic received little attention in the past (exceptions are [23, 24, 20, 22]).

Thirdly, the way in which the HT was adapted to extract line segments, *i.e.*, by taking the output of the HT and obtaining the start and end points of the segments, originates issues of its own. While initially each point in the Hough space accumulated votes supporting the existence of only one image line, now the votes can refer to multiple collinear line segments of various lengths and different start and end points. Since only the number of collinear edge points is stored in a conventional Hough space, the votes of individual line segments cannot be distinguished. This means that a local maximum in the Hough space no longer implies a maximum likelihood that a line or line segments actually exist in the image with such parameterization — therefore, this adaptation of the HT is not a statistically robust line segment extractor.

Few papers deal with this topic as well, as in [20, 22].

Another issue of the HT is that it cannot deal properly with wide line segments [24], *i.e.*, segments made up of blurred edges, since the highest number of votes corresponds to the diagonal of the segment rather than the segment itself. The limitations of the HT in handling complex images have been pointed out by several authors, *e.g.*, [21, 9, 22], and many efforts have been made to alleviate its problems. They include, *e.g.*, the use of the edge direction to reduce the accumulation of spurious votes in the Hough space [25, 26], the sequential processing and removal of the strongest peaks in the Hough space [21, 16], sub-sampling of the edge map (randomized HT) [27]. Other authors addressed storage and computational issues of the HT by proposing a hierarchical scheme [28], multiple accumulator resolutions [29] and a probabilistic formulation [30]. The thickness of line segments and edge point connectivity is used in [24] to change the value of each vote in a standard HT. However, none of these methods tackle the fundamental issues of the HT in extracting line segments in complex images.

The other set of popular methods for line segment extraction can be categorized as *local methods*, due to their reliance on local decisions rather than global ones (see [31, 32, 33, 34, 9] for examples). The majority of local methods use three steps: first, they detect edge points and chain them (using, *e.g.*, the method in [35]); then, a rough estimate of the segment direction is computed using, *e.g.*, total least-squares regression [34]; and finally, they refine and extend the segment by including new edge points that approximately fit the line. The final step usually involves alternating between two stages until convergence [34]: inclusion of new edge points that are close to the candidate line, according to a distance measure; and re-estimation of the line segment parameters from the new set of edge points. In realistic scenarios, the initial step of detecting edge points and large connected regions belonging to a single segment is hard due to texture, low-contrast regions, crossing segments, and noise. The resulting smaller edge point chains hurt the reliability of the regression step and, finally, as it is typical with this type of alternating methods, a poor initial model for the line segment model may compromise the final refinement and extension results. Variations of the basic method include, *e.g.*, [36], which takes the chained edge points and cuts them into line segments, using a straightness criterion. References [33, 37] and [38] bypass the chaining of edge points and fit a line directly to all edge points inside a sliding window. The segment direction is estimated roughly in [39] using a so-called local HT and taking the peaks of local orientation

histograms, computed at each edge point. Two popular local methods for line segment detection are [32] and the LSD (Line Segment Detector) of [9]. The method in [32] coarsely quantizes the local orientation angles, chains adjacent pixels with identical orientation labels, and fits a line segment to the grouped pixels. LSD [9] extends this idea by using continuous angles and eliminates false line segment detections by using the Helmholtz principle of [40]. These methods result computationally simple but lack robustness to deal with the imperfections that occur in realistic scenarios.

In [22], we propose a better adaptation of the HT to the extraction of line segments, which retains its global aspect and solves its main issues in dealing with complex images. It starts by computing image derivatives using small filters and obtaining the directions for which there is a predominance of positive or negative derivatives (corresponding to a dark-to-bright transition in the image or vice-versa). Such predominance gives a rough indication that a line segment might be present at that location (it implements the sign test, as we discuss later) and only edge points within areas of strong predominance are allowed to contribute. Then, instead of using a single Hough space, where collinear line segments are indiscernible, we use a local HT for each line segment. Connectivity is incorporated in the voting process, by only accounting for the contributions of edge points whose position and directional content agree with potential line segments. As a consequence, the vast majority of spurious votes are eliminated and the peaks in the local Hough spaces correspond to line segments of maximum length. Unfortunately, the computational complexity of the proposed method is prohibitive for many applications (computation times of 100 to 1000 seconds are reported in [22]) and thick transitions are not obtained correctly, as in the standard HT. We now believe that the requirements of global methods, of storing comprehensive data in order to eliminate early decisions that compromise robustness, make such methods computationally too complex for most applications.

1.2. Proposed approach

The novel and key aspect of our approach is the combination of what we denote *contextual* and *local edges*. Contextual edges are thresholded image derivatives obtained with filters of large footprint, which reduce the influence of noise when identifying image transitions. Although large footprint filters are robust to noise, edge localization is imprecise since every transition is smeared by their large point spread function. On the other hand, local edges are thresholded image derivatives obtained with filters of very small footprint,

e.g., Sobel, Prewitt, and Roberts operators. Since these filters are small, edges are localized precisely but noise may originate erroneous detections of multiple directions and amplitudes. Our proposal combines contextual and local edges obtained at the same pixels by taking the sign of the contextual edge and using it to identify so-called valid local edges with the same sign. The edge sign indicates if it is a dark–to–bright transition or the opposite.

In complex images, valid local edges are disconnected with each other, due to noise and clutter. To obtain connected edge maps, we handle connectivity explicitly in the combination process by checking if the valid local edges, along a given direction, are at a distance not greater than a maximum distance threshold from each other. If so, the pixels corresponding to valid local edges and those in between are marked as edge points. The resulting edge detector has the robustness of contextual edges in dealing with noise and the localization of local edges, as idealized by Canny [18]. Since the combination of contextual and local edges originates pixel-thin (connected) regions of the length of the line segment, a simple region growing and rectangle fitting methodology (which we detail later) suffices for the individual thin segments to be combined into line segments of all lengths and widths.

Typical contextual edge detectors handle noise by applying a low-pass filter of large footprint such as, *e.g.*, Gabor or steerable filters [41], followed by a derivation step and binarization with a fixed threshold. However, since other unwanted high-frequency variations due to, *e.g.*, textures and clutter from interfering line segments and non-rectilinear image data, are common in complex images and have unpredictable amplitudes — unlike typical image white noise, which is constant throughout the image — an adaptive threshold is beneficial. For this purpose, we take both the mean and the variance of the two sets of pixels and use a two-sample statistical test to determine if a contextual edge exists. We consider that each set of pixels follows a Normal distribution with the sampled parameters, compute the Total Variation distance between them and threshold the confidence interval for the null hypothesis that both distributions are actually the same.

By using large (yet limited) contextual filters, our semi-global method results simple. Also, by doing away with the issues of local edges in a more effective way than typical global methods, through a combination with contextual edges, it results typically more robust as well in obtaining line segments of all lengths and transition widths. This is demonstrated in the experimental section, where we present a complexity analysis and illustrative results using synthetic and real images to compare our method with other

methods: the standard HT [13], the state-of-the-art of local methods LSD [9], and our previous HT-based method [22].

1.3. Paper organization

The organization of the remaining of the paper is as follows. In section 2 we provide a brief introduction to statistical edge detectors and detail our implementation. Section 3 details the combination of contextual and local edges, including the explicit tackling of connectivity. Rectangle fitting is described in 4 and the experimental results are reported in Section 5. Section 6 concludes the paper.

2. Statistical edge detection

Low-pass filtering prior to derivation is used extensively in edge detection schemes to reduce the effect of noise, albeit with small footprints for small localization error. This can be seen roughly as computing the mean of two sets of M pixels, comparing them and binarizing the result using threshold C. In this scheme, the value of C is critical, since a small C originates erroneous edges, and a large C misses them. In simpler scenarios where noise has a constant variance σ^2 throughout the image, the optimal threshold between two sets of pixels can be defined optimally for a given confidence interval (*e.g.*, the standard error of the mean, $C = \sigma/\sqrt{M}$). However, since complex images contain unwanted high-frequency variations of unpredictable variance such as, *e.g.*, textures and clutter from other image data, a fixed threshold is far from optimal. This suggests that a more comprehensive statistical analysis is able to improve edge detection. Such studies were initiated by Bovik et al. [42], which used two-sample statistical tests in the context of edge detection. Two-sample tests take two sets of pixel values, illustrated in Fig. 1, and considers that they are samples of two underlying distributions, $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$. The test then checks the null hypothesis $\mathcal{H}_0$ that the underlying probability distributions are in fact the same [43]. A contextual edge exists between the samples of $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$ if the distributions are deemed different.

2.1. Typical approaches

The parameters of each distribution can be estimated from the samples by computing $\hat{\mu}_i = \frac{1}{M}\sum_{j=1}^{M} \boldsymbol{x}_{ij}$ and $\hat{\sigma}_i^2 = \frac{1}{M-1}\sum_{j=1}^{M} \left(\boldsymbol{x}_{ij} - \hat{\mu}_i\right)^2$ (with $i \in \{T, B\}$). If the two sets of samples are assumed to be Normally distributed, $\boldsymbol{X}_T \sim \mathcal{N}(\mu_T, \sigma_T^2)$ and $\boldsymbol{X}_B \sim \mathcal{N}(\mu_B, \sigma_B^2)$, various criteria can then be used.

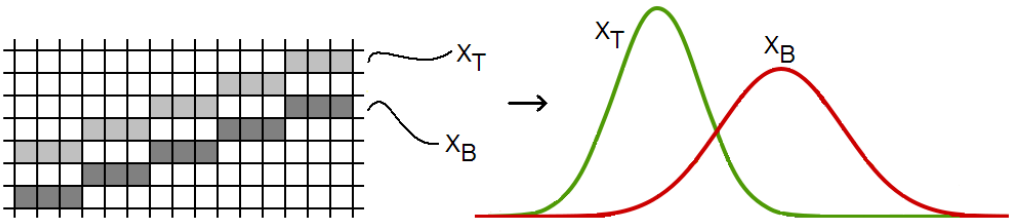

Figure 1: Illustration of two-sample tests. The pixels (left) are samples of underlying distributions $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$ (right). The two-sample test determines if they are the same.

For the t-test, the distributions are the same if their means coincide. Since the sample mean varies with the sample variance and the number of samples through formula $\sigma_{\hat{\mu}_i} = \hat{\sigma}_i/\sqrt{M}$, the t-test is given by

$$t = (\hat{\mu}_T - \hat{\mu}_B) \Big/ \sqrt{\frac{\hat{\sigma}_T^2 + \hat{\sigma}_B^2}{M}} \text{ [43]}. \tag{1}$$

Once a t-value is determined, the probability that the test statistic would take a value at least as extreme as the one observed, denoted p-value, is computed, for $M-1$ degrees of freedom. If the calculated p-value is below the threshold chosen for statistical significance (usually the 0.10, the 0.05, or 0.01 level), the distributions are deemed different (two-tailed test) or larger than the other (one-tailed test) [43].

If no assumption is made about the distributions of $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$, non-parametric tests can be used instead. These tests are more general and robust to outliers [43] but also computationally more intensive, since they require expensive sorting operations. Some tests compute empirical cumulative distributions (ecd) from samples $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$ and then compare them using some distance measure. The Kolmogorov–Smirnov test obtains the maximum distance between the ecds and the Fisz–Cramér–Von Mises test integrates the squared difference between the ecds. The Wilcoxon Mann–Whitney is a popular rank order test that mixes the samples of both distributions, sorts and ranks them. The difference between the distributions is assessed by adding the ranks of one distribution and comparing with the added ranks of the other [43].

The success of two-sample statistical tests in contextual edge detection in noisy images is shown in [44]. Reference [45] uses a t-test for detecting edges and a mixture of Normal distributions models noisy data in the edge detector of [46]. Various statistical tests are used in the neural network approach [47].

In our previous line segment extraction method [22], the number of local positive image derivatives minus the number of negative ones is used to assess the predominance of dark-to-bright transitions or vice-versa, indicating a high chance of having found a line segment. This procedure implements the sign test, a nonparametric paired test that tests the null hypothesis that there is "no difference in medians" between $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$. Despite lacking the statistical power of other tests, it has very general applicability, as it makes very few assumptions about the nature of the distributions under test [43].

2.2. Our approach

Since nonparametric tests are computationally expensive, we use a parametric test instead. Although the t-test is arguably the most frequently used parametric test, it assumes that coinciding mean values imply identical distributions (this is clear in equation (1)), which is not coherent with the intuition that pixel distributions with the same sample mean but different sample variances can be deemed visually different. To enable this feature, we use the Total Variation (TV) distance between two Normal distributions [48],

$$\delta(\boldsymbol{X}_T, \boldsymbol{X}_B) = \frac{1}{2}\int_{-\infty}^{\infty} \left| f(\xi; \hat{\mu}_T, \hat{\sigma}_T^2) - f(\xi; \hat{\mu}_B, \hat{\sigma}_B^2) \right| d\xi \in [0, 1], \qquad (2)$$

where $f(\cdot; \mu, \sigma^2)$ represents the probability density function (pdf) of the Normal distribution. The TV distance outputs the integral of the (linear) distance between the pdfs of $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$. As we show below, the linear aspect assures that only the ratio between the sample standard deviations is taken into account, instead of their actual values. Although more comprehensive tests are needed, the higher sensitivity and robustness of the TV distance in detecting the perceived line segments, in contrast with non-linear distances such as the Kullback-Leibler and Hellinger divergences, lead us to believe that this feature of the TV distance best emulates the human visual system in estimating the boundary strength between different areas.

To simplify the calculation of (2), let $\xi_i \in \{-\infty, \hat{\xi}_1, \hat{\xi}_2, \infty\}$ be the points in which the pdf of distributions $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$ are equal, as illustrated in Fig. 2. Points ξ_i (with $\xi_i \leq \xi_{i+1}$, $\forall i \in \{1, \ldots, 3\}$) determine where one pdf becomes larger than the other and, thus, helps in dealing with the magnitude operator and enables the use of the cumulative density function of the Normal

distribution, $F(\cdot;\mu,\sigma^2)$,

$$\begin{aligned}\delta(\boldsymbol{X}_T,\boldsymbol{X}_B) &= \frac{1}{2}\sum_{i=1}^{3}\int_{\xi_i}^{\xi_{i+1}} \left|f(\xi;\hat{\mu}_T,\hat{\sigma}_T^2) - f(\xi;\hat{\mu}_B,\hat{\sigma}_B^2)\right| d\xi \\ &= \frac{1}{2}\sum_{i=1}^{3}\left|\int_{\xi_i}^{\xi_{i+1}} f(\xi;\hat{\mu}_T,\hat{\sigma}_T^2) - f(\xi;\hat{\mu}_B,\hat{\sigma}_B^2) d\xi\right| \qquad (3) \\ &= \frac{1}{2}\sum_{i=1}^{3}\left|F(\xi_{i+1};\hat{\mu}_T,\hat{\sigma}_T^2) - F(\xi_i;\hat{\mu}_T,\hat{\sigma}_T^2) - F(\xi_{i+1};\hat{\mu}_B,\hat{\sigma}_B^2) + F(\xi_i;\hat{\mu}_B,\hat{\sigma}_B^2)\right|,\end{aligned}$$

To determine $\hat{\xi}_1$ and $\hat{\xi}_2$, we make $f(\hat{\xi};\hat{\mu}_T,\hat{\sigma}_T^2) = f(\hat{\xi};\hat{\mu}_B,\hat{\sigma}_B^2)$, which results in $\hat{\xi} = \left(-b \pm \sqrt{b^2-4ac}\right)/2a$, where $a = 1/\left(2\hat{\sigma}_T^2\right) - 1/\left(2\hat{\sigma}_B^2\right)$, $b = -\hat{\mu}_T/\left(\hat{\sigma}_T^2\right) + \hat{\mu}_B/\left(\hat{\sigma}_B^2\right)$, and $c = \hat{\mu}_T/\left(2\hat{\sigma}_T^2\right) - \hat{\mu}_B/\left(2\hat{\sigma}_B^2\right) - \ln\left(\hat{\sigma}_B/\hat{\sigma}_T\right)$. If all parameters are identical, then $\delta(\boldsymbol{X}_T,\boldsymbol{X}_B) = 0$. If only the sample variances are identical, the pdf of both distributions cross at only one point, $\hat{\xi}_1 = \hat{\xi}_2 = \left(\hat{\mu}_T - \hat{\mu}_B\right)/2$. For different sample variances, two crossing points are obtained.

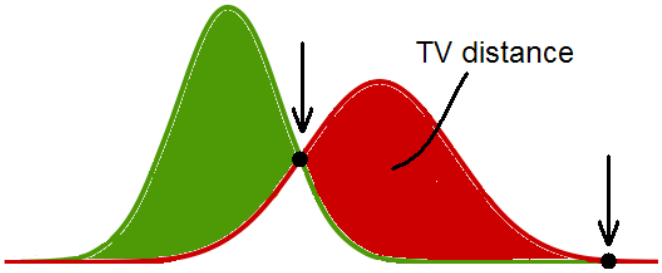

Figure 2: Illustration of distribution pdf intersection points and area for TV distance.

To eliminate redundant calculations in the remainder of our method, we compute a look-up table with relevant TV distances. Since the TV distance requires four parameters, $(\hat{\mu}_T,\hat{\sigma}_T,\hat{\mu}_B,\hat{\sigma}_B)$, it seems that a four-dimensional look-up table is needed, which is expensive to compute and store. However, since we are computing the linear distance between two distributions, it is not the actual values of $\hat{\mu}_T$ and $\hat{\mu}_B$ that are needed but how they relate with each other. Similarly, the actual values of $\hat{\sigma}_T$ and $\hat{\sigma}_B$ are not needed but only the relation between them as well. Let variables $\hat{\mu}'_B$ and $\hat{\sigma}'_B$ be

$$\hat{\mu}'_B = \frac{|\hat{\mu}_B - \hat{\mu}_T|}{\min(\hat{\sigma}_T,\hat{\sigma}_B)} \quad \hat{\sigma}'_B = \frac{\max(\hat{\sigma}_T,\hat{\sigma}_B)}{\min(\hat{\sigma}_T,\hat{\sigma}_B)} \ . \qquad (4)$$

Assuming that $\hat{\mu}_T \leq \hat{\mu}_B$ and $\hat{\sigma}_T \leq \hat{\sigma}_B$, equations (4) become

$$\hat{\mu}_B = \hat{\mu}_T + \hat{\mu}'_B \hat{\sigma}_T \quad \hat{\sigma}_B = \hat{\sigma}_T \hat{\sigma}'_B. \tag{5}$$

Replacing (5) in (2) and substituting variable $\zeta = (\xi - \mu_T)/\sigma_T$, we have

$$\begin{aligned} \delta(\boldsymbol{X}_T, \boldsymbol{X}_B) &= \frac{1}{2} \int_{-\infty}^{\infty} \left| f(\xi; \hat{\mu}_T, \hat{\sigma}_T^2) - f(\xi; \hat{\mu}_T + \hat{\mu}'_B \hat{\sigma}_T, \hat{\sigma}_T^2 \hat{\sigma}_B'^2) \right| d\xi \\ &= \frac{1}{2} \int_{-\infty}^{\infty} \left| f(\zeta; 0, 1)/\hat{\sigma}_T - f(\zeta; \hat{\mu}'_B, \hat{\sigma}_B'^2)/\hat{\sigma}_T \right| \hat{\sigma}_T d\zeta \\ &= \frac{1}{2} \int_{-\infty}^{\infty} \left| f(\zeta; 0, 1) - f(\zeta; \hat{\mu}'_B, \hat{\sigma}_B'^2) \right| d\zeta, \end{aligned} \tag{6}$$

which shows that the TV distance depends only on two parameters (analogous demonstration show the same result in any scenario). A two-dimensional look-up table $\boldsymbol{\Delta}(\hat{\mu}'_B, \hat{\sigma}'_B)$ is then filled in for various $(\hat{\mu}'_B, \hat{\sigma}'_B)$, with $\hat{\mu}_B \geq 0$ and $\hat{\sigma}'_B \geq 1$, using (3). During normal operation, the parameters are normalized using equations (4) and the corresponding TV distance is accessed.

3. Combining contextual and local edges using connectivity

In this section, we compute contextual edges and combine them with local ones to form connected edge maps. Since samples $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$, the inputs of the two-sample test at the heart of contextual edge extraction, are taken from long and thin footprints (as illustrated in Fig. 1), each test is effective only in a small angular range. Therefore, contextual edges must be computed along multiple directions to span the entire 180° range. For every direction, we use a running average approach to take samples $\boldsymbol{X}$ at each point and, assuming that they follow a Normal distribution, obtain their parameters. Then, we take the parameters refering to $\boldsymbol{X}_T$ and $\boldsymbol{X}_B$, compute the Total Variation distance detailed in the previous chapter and describe each pixel (and direction) as containing a *negative*, a *positive* or *no contextual edge*, using quantization thresholds defined later. If there is a contextual edge, it is used to select *valid* local edges. Local edges consist of image derivatives obtained by convolving the image with kernels of very small footprint, central difference kernels, whose results are also described as being a negative, a positive or no local edge. Valid local edges then consist of local edges that have the same sign as the contextual one. If the distance between valid local edges does not exceed a maximum distance threshold d, those edges and the pixels between them are marked as edge points.

3.1. Computing Normal distribution parameters

The parameters of the Normal distribution are computed for each set of samples starting at pixel $\boldsymbol{p}_m$, illustrated in Fig. 3. We use $M = 15$ pixels at each set of samples and we verified theoretical and experimentally that $N = 32$ different uniformly spaced directions leads simultaneously to neglectable angular coverage errors and good computational performance (since every pixel in the perimeter of a semicircular window of radius M is used only when $N \geq \pi(M-1)$, $N = 44$ different directions should be used for $M = 15$ pixels and perfect angular performance).

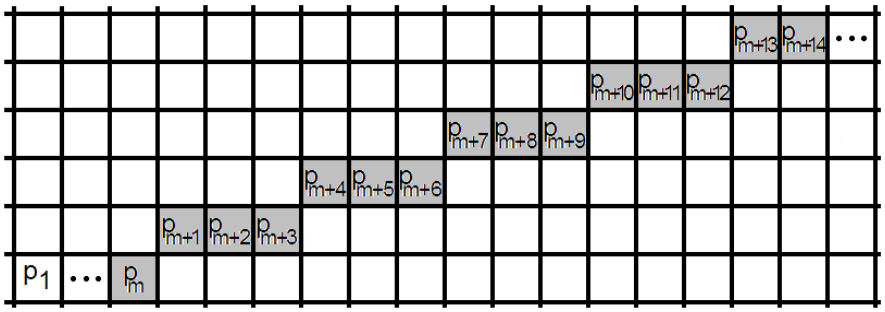

Figure 3: Illustration of $M = 15$ pixels starting at pixel $\boldsymbol{p}_m$ along direction θ_n.

Each direction $n \in \{1, \ldots, N\}$ corresponds to angle $\theta_n = 180°(n-1)/N$, which lies in the horizontal (H) or vertical (V) half of the semicircle $H(\theta_n)$, and in quadrant $Q(\theta_n)$ of the semicircle, both illustrated in Fig. 4,

$$H(\theta_n) = \begin{cases} \text{V} & \theta_n \in [45°, 135°[\\ \text{H} & \textit{otherwise} \end{cases}, \tag{7}$$

$$Q(\theta_n) = \begin{cases} 0 & \theta_n \in [0°, 22.5°[\cup [157.5°, 180°[\\ 45 & \theta_n \in [22.5°, 67.5°[\\ 90 & \theta_n \in [67.5°, 112.5°[\\ 135 & \theta_n \in [112.5°, 157.5°[\end{cases}.$$

Our method works by selecting a direction θ_n and dividing the image into lines along that direction, as illustrated in Fig. 3. For each point $\boldsymbol{p}_m$ in a line, the parameters of the Normal distribution are computed by taking M consecutive pixels, from pixel $\boldsymbol{p}_m$ to pixel $\boldsymbol{p}_{m+M-1}$, and computing their sample average and variance, $\hat{\boldsymbol{\mu}}(\boldsymbol{p}_m)$ and $\hat{\boldsymbol{\sigma}}^2(\boldsymbol{p}_m)$. Since the set of M consecutive pixels needed for the next point in the line, $\boldsymbol{p}_{m+1}$, is the same as

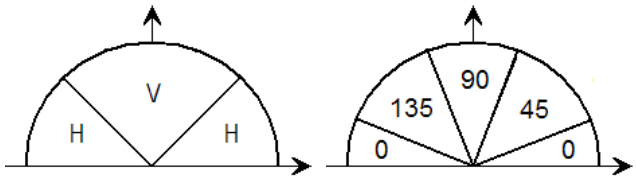

Figure 4: Illustration of the half of a semicircle, $H(\theta_n)$ (left), and quadrant, $Q(\theta_n)$ (right).

the set of points for $\boldsymbol{p}_m$ except for the points at the start and end of the set, we use a recursion that simplifies calculations and is used often: the running average. To formalize this idea, for image $\boldsymbol{I} \in \mathbb{R}^{S_x \times S_y}$, we define a mapping that addresses each line along direction θ_n,

$$\boldsymbol{A}(x, y, \theta_n) = \begin{cases} (x, y + [x \tan \theta_n] + \gamma S_y) & \text{if } H(\theta_n) = H \\ (x + [y \tan^{-1} \theta_n] + \gamma S_x, y) & \text{if } H(\theta_n) = V \end{cases}, \tag{8}$$

where $[\cdot]$ refers to the rounding operation and integer parameter γ is chosen so that $\boldsymbol{A}(x, y, \theta_n)$ lies inside the image limits. If $H(\theta_n) = H$, the y-th line is made up of the pixels addressed by $\boldsymbol{A}(x, y, \theta_n)$, with $x \in \{1, \ldots, S_x\}$ and in an increasing order. Analogously, if $H(\theta_n) = V$, the x-th line is addressed by $\boldsymbol{A}(x, y, \theta_n)$, with increasing $y \in \{1, \ldots, S_y\}$. To compute the average and standard deviation in a recursive way, we define linear and quadratic accumulators, Φ_x and Φ_{x^2}, respectively, which are updated using

$$\begin{aligned} \Phi_x &\leftarrow \Phi_x - \boldsymbol{I}(\boldsymbol{p}_{m-1}) + \boldsymbol{I}(\boldsymbol{p}_{m+M-1}) \\ \Phi_{x^2} &\leftarrow \Phi_{x^2} - \boldsymbol{I}^2(\boldsymbol{p}_{m-1}) + \boldsymbol{I}^2(\boldsymbol{p}_{m+M-1}), \end{aligned} \tag{9}$$

where $\boldsymbol{p}_m$ locates the current pixel, $\boldsymbol{p}_{m-1}$ locates the pixel that is exiting the set of M pixels, and $\boldsymbol{p}_{m+M-1}$ locates the pixel that is entering the set,

$$\begin{aligned} &\boldsymbol{p}_m = \boldsymbol{A}(x, y, \theta_n), \\ &\boldsymbol{p}_{m-1} = \boldsymbol{A}(x-1, y, \theta_n), \quad \boldsymbol{p}_{m+M-1} = \boldsymbol{A}(x+M-1, y, \theta_n) \quad \text{if } H(\theta_n) = H \\ &\boldsymbol{p}_{m-1} = \boldsymbol{A}(x, y-1, \theta_n), \quad \boldsymbol{p}_{m+M-1} = \boldsymbol{A}(x, y+M-1, \theta_n) \quad \text{if } H(\theta_n) = V \end{aligned} . \tag{10}$$

The parameters of the Normal distribution, for direction θ_n, are then

$$\begin{aligned} \hat{\boldsymbol{\mu}}(\boldsymbol{p}_m) &= \frac{1}{M} \Phi_x \\ \hat{\boldsymbol{\sigma}}^2(\boldsymbol{p}_m) &= \frac{1}{M-1} \Phi_{x^2} - \frac{1}{M(M-1)} \Phi_x^2, \end{aligned} \tag{11}$$

for each point in the image.

3.2. Combining contextual and local edges using connectivity

To compute local edges, we first capture the local directional content of image $\boldsymbol{I}$ by computing its derivatives, through the convolution with four oriented kernels,

$$\boldsymbol{\nabla}_\theta \boldsymbol{I} = \boldsymbol{I} * \boldsymbol{K}_\theta \,, \quad \theta \in \{0°, 45°, 90°, 135°\} \,. \tag{12}$$

Local orientation has been exploited before and captured by using several types of kernels, see, *e.g.*, [32, 25, 9]. Although kernels with a large support would smooth the noise, we use the same approach as in our previous work [22] and employ simple central difference kernels, since they enable more precise edge localization, by minimizing the influence of surrounding pixels. Thus, we set

$$\boldsymbol{K}_0 = \begin{bmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & -1 & 0 \end{bmatrix}, \boldsymbol{K}_{45} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -1 \end{bmatrix}, \boldsymbol{K}_{90} = \begin{bmatrix} 0 & 0 & 0 \\ 1 & 0 & -1 \\ 0 & 0 & 0 \end{bmatrix}, \boldsymbol{K}_{135} = \begin{bmatrix} 0 & 0 & -1 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{bmatrix}. \tag{13}$$

The combination of contextual and local edges occurs in three steps and is described in detail in the sequel. To summarize, in the first step, the parameters of the Normal distributions, for all pixels and directions (computed in 3.1), and the local image derivatives (computed above) are used to compute contextual and local edges at each pixel. This occurs for each direction θ_n and progressively along each line, and the goal is to find a contextual and local edge with matching sign. Once both edge types are matched, at pixel $\boldsymbol{p}_m$, a second step then checks if the M pixels between $\boldsymbol{p}_m$ and $\boldsymbol{p}_{m+M-1}$ contain valid local edges that are sufficiently connected with each other, as should occur in a line segment. We consider that valid local edges are connected if they are not separated by more than a maximum distance threshold d. If step 2 is successful, the beginning of a line segment was found. A third step marks those pixels as *connected edge points* and progressively checks and marks the subsequent pixels along direction θ_n which preserve the contextual edge sign and whose valid local edges continue to be sufficiently connected. The output of this process is a set of connected edge points for each direction θ_n, which is used in section 4 to fit rectangles.

3.2.1. Search for initial edge

In this step, the pixels along direction θ_n are scanned to find a contextual and local edge whose spatial location and sign coincides. For $H(\theta_n) = H$

(the adaptation to $H(\theta_n) = V$ is analogous), each y-th line is scanned at a time ($y \in \{1, \ldots, S_y\}$) and, within each line, the pixels are scanned for all $x \in \{1, \ldots, S_x\}$ in increasing order, resulting in coordinate $\boldsymbol{p}_1 = \boldsymbol{A}(x, y, \theta_n)$ (to simplify the explanation of our method, we define that $\boldsymbol{p}_m$ refers to the m-th pixel of the set of M pixels currently being tested).

To determine if there is a contextual edge at point $\boldsymbol{p}_1$ along direction θ_n, we start by computing the coordinates of the parameters above and below $\boldsymbol{p}_1$ (see Fig. 5 for an illustration), $\boldsymbol{p}_{1_T} = \boldsymbol{A}(x - [\sin\theta_n], y - [\cos\theta_n], \theta_n)$ and $\boldsymbol{p}_{1_B} = \boldsymbol{A}(x + [\sin\theta_n], y + [\cos\theta_n], \theta_n)$, and obtain the parameters of the two distributions, $\hat{\mu}_T = \hat{\boldsymbol{\mu}}(\boldsymbol{p}_{1_T})$, $\hat{\sigma}_T^2 = \hat{\boldsymbol{\sigma}}^2(\boldsymbol{p}_{1_T})$, $\hat{\mu}_B = \hat{\boldsymbol{\mu}}(\boldsymbol{p}_{1_B})$, $\hat{\sigma}_B^2 = \hat{\boldsymbol{\sigma}}^2(\boldsymbol{p}_{1_B})$. The parameters are normalized using equation (4) and the two-dimensional look-up table containing the TV distances is accessed, $\boldsymbol{\Delta}(\mu'_B, \sigma'_B)$. Since the type of transition is important for line segment extraction, *i.e.*, if it is a dark-to-light transition or vice-versa, contextual and local edges should include this information. For this effect, the contextual edge is given by $\delta_{X_T X_B} = \boldsymbol{\Delta}(\mu'_B, \sigma'_B)\text{sgn}(\hat{\mu}_T - \hat{\mu}_B)$. A contextual edge exists if $|\delta_{X_T X_B}| \geq C$.

There is a local edge at pixel $\boldsymbol{p}_1$ if $|\boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_1)| \geq L_C$. If it has the same sign as the contextual one (*i.e.*, if condition $\text{sgn}(\delta_{X_T X_B}) = \text{sgn}\left(\boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_1)\right)$ is true), it is denoted as a *valid* local edge. If a valid local edge was found, the method progresses onto step 2. Threshold L_C is given by $L_C = \max(L, |\hat{\mu}_T - \hat{\mu}_B|/2)$, where $L = 3$ and $C = 0.7$. The overall combination of contextual and local edges using connectivity is summarized in Alg. 1, for $H(\theta_n) = H$.

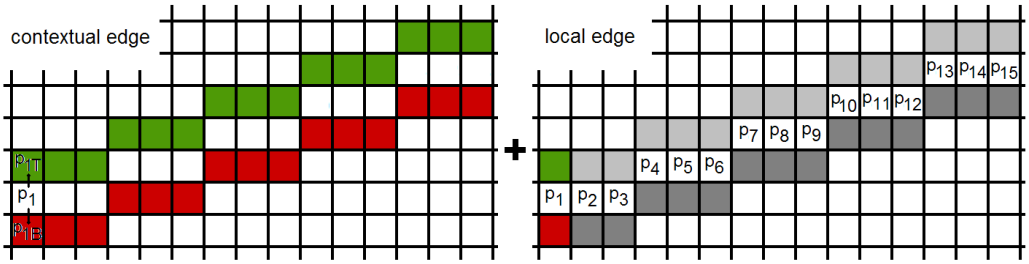

Figure 5: Search for initial contextual (left) and valid local edge (right).

3.2.2. Marking of initial area

The contextual and valid local edge found in step 1 indicates that a set of connected edge points may start at that location. That occurs if the M pixels between the top and bottom set of samples contain valid local edges

Algorithm 1 combination of contextual and local edges using connectivity

input: directional data $\boldsymbol{\nabla}_\theta \boldsymbol{I}$, parameters $\hat{\boldsymbol{\mu}}$ and $\hat{\boldsymbol{\sigma}}^2$, maximum distance threshold d, contextual C and local L thresholds, TV distances $\boldsymbol{\Delta}(\cdot,\cdot)$
$\boldsymbol{E} \leftarrow 0$
% for every pixel in the image (assuming $H(\theta_n) = H$)
for $y \in \{1, \ldots, S_y\}$ **do**
  step = 'search for start'
  for $x \in \{1, \ldots, S_x\}$ **do**
    % get contextual edge data
    $\boldsymbol{p}_{1_T} = \boldsymbol{A}(x - [\sin\theta_n], y - [\cos\theta_n], \theta_n)$
    $\boldsymbol{p}_{1_B} = \boldsymbol{A}(x + [\sin\theta_n], y + [\cos\theta_n], \theta_n)$
    $\hat{\mu}_T = \hat{\boldsymbol{\mu}}(\boldsymbol{p}_{1_T})$, $\hat{\sigma}_T^2 = \hat{\boldsymbol{\sigma}}^2(\boldsymbol{p}_{1_T})$, $\hat{\mu}_B = \hat{\boldsymbol{\mu}}(\boldsymbol{p}_{1_B})$, $\hat{\sigma}_B^2 = \hat{\boldsymbol{\sigma}}^2(\boldsymbol{p}_{1_B})$
    $\mu'_B = |\hat{\mu}_T - \hat{\mu}_B| / \min(\hat{\sigma}_T, \hat{\sigma}_B)$, $\sigma'_B = \max(\hat{\sigma}_T, \hat{\sigma}_B) / \min(\hat{\sigma}_T, \hat{\sigma}_B)$
    $\delta_{X_T X_B} = \boldsymbol{\Delta}(\mu'_B, \sigma'_B)\mathrm{sgn}(\hat{\mu}_T - \hat{\mu}_B)$
    $L_C = \max(L, |\hat{\mu}_T - \hat{\mu}_B|/2)$
    % combine contextual and local edges
    if step = search for start' **then** Alg. 2 **else** Alg. 3 **end if**
  end for
end for
output: Edge map $\boldsymbol{E}$

that are sufficiently connected to each other. Step 2 verifies this and, if true, marks the pixels as *connected edge points*.

Pixels $\boldsymbol{p}_m \in \{\boldsymbol{p}_1, \ldots, \boldsymbol{p}_M\}$ are checked sequentially, as illustrated in Fig. 6 (where $\boldsymbol{p}_m = \boldsymbol{A}(x + m - 1, y, \theta_n)$, if $H(\theta_n) = H$). Whenever a valid local edge is found, *i.e.*, $\mathrm{sgn}\,(\delta_{X_T X_B})\, \boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_m) \geq L_C$, gap counter g_C is nulled — otherwise, g_C is incremented. Note that, while in approaches such as [23], the distance between binary local edge points is the only criteria to judge whether pixels are connected with each other or not, our approach (as also illustrated in Fig. 6) requires that the sign of the edge points matches the sign of the connected edges. Local edges with opposite signs — even strong ones — are discarded in the same way as non-edges.

If the gap counter exceeds the maximum distance threshold d, the gap between two valid edge points is too large and step 2 ends unsuccessfully, returning to step 1 for further searching. If the gap counter never exceeded d, the pixels are marked as edge points of sign $\boldsymbol{E}(\boldsymbol{p}_m) = \mathrm{sgn}\,(\delta_{X_T X_B})$ and the method proceeds to step 3. We allow gaps of up to $d = 5$ pixels in this

paper. Step 1 and 2 of the combination of contextual and local edges are summarized in Alg. 2, for $H(\theta_n) = H$.

Algorithm 2 find initial set of connected edge points (assumes $H(\theta_n) = H$)

input: signed TV distance $\delta_{X_T X_B}$, directional data $\boldsymbol{\nabla}_{Q(\theta_n)}\boldsymbol{I}$, maximum distance threshold d, contextual C and local L_C thresholds
$\boldsymbol{p}_1 = \boldsymbol{A}(x, y, \theta_n)$
if $|\delta_{X_T X_B}| \geq C$, $|\boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_1)| \geq L_C$, $\delta_{X_T X_B}\boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_1) > 0$ **then**
  % found initial edge, implement step 2
  while $m \in \{1, \ldots, M\}$ and $g_C \leq d$ **do**
    if $\operatorname{sgn}(\delta_{X_T X_B})\,\boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_m = \boldsymbol{A}(x + m - 1, y, \theta_n)) \geq L_C$ **then**
      % found valid local edge
      $g_C \leftarrow 0$
    else
      $g_C \leftarrow g_C + 1$
    end if
  end while
  if $g_C \leq d$ **then**
    % does not exceed d, mark M edge points
    for $m \in \{1, \ldots, M\}$ **do**
      $\boldsymbol{E}(\boldsymbol{p}_m = \boldsymbol{A}(x + m - 1, y, \theta_n)) = \operatorname{sgn}(\delta_{X_T X_B})$
    end for
    step = 'mark area'
  end if
end if

3.2.3. Marking of connected area

As the set of connected edge points increase, the contextual and local edges to be checked move progressively along the line. For $H(\theta_n) = H$, this occurs by making $x \leftarrow x + 1$ and re-computing both types of edges for contextual and local compliance. The contextual edge is computed as before, by accessing the parameters at positions $\boldsymbol{p}_{1_T}$ and $\boldsymbol{p}_{1_B}$ and computing the new value of $\delta_{X_T X_B} = \boldsymbol{\Delta}(\mu'_B, \sigma'_B)\text{sgn}(\hat{\mu}_T - \hat{\mu}_B)$. If $|\delta_{X_T X_B}| \geq C$ and has the sign of step 1, the contextual edge is valid.

To determine if there a valid local edge, pixel $\boldsymbol{p}_M = \boldsymbol{A}(x + M - 1, y, \theta_n)$ is checked. If $\operatorname{sgn}(\delta_{X_T X_B})\,\boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_M) \geq L_C$, it a valid local edge and counter

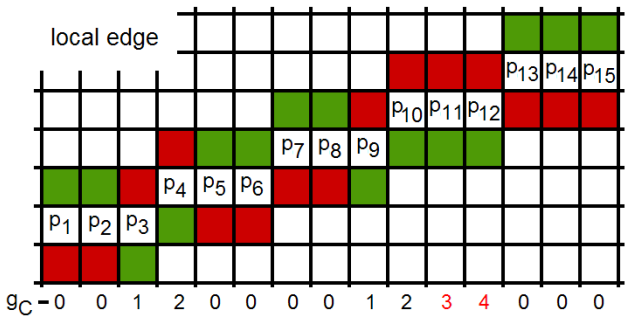

Figure 6: Illustration of the sign of local edges and gap counter, g_C (using the contextual edge illustrated in Fig. 5). Using $d = 2$ (for illustration), the four consecutive local edges of contrary sign makes step 2 end abruptly while analyzing pixel $\boldsymbol{p}_{11}$.

g_C is nulled — otherwise, g_C is incremented. If $g_C \leq d$ and the contextual edge is valid, pixel $\boldsymbol{p}_M$ is marked as a connected edge point of sign $\boldsymbol{E}(\boldsymbol{p}_M) = \operatorname{sgn}(\delta_{X_T X_B})$. Otherwise, the set of connected edge points ended. The last g_C non-connected edge points are unmarked and the method returns to step 1, to search for a new set. This is illustrated in Fig. 7 and summarized in Alg. 3.

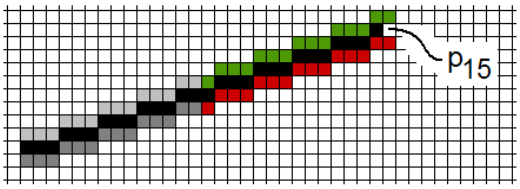

Figure 7: Illustration of the marking of connected edge points. The connected edge points are marked in black. Contextual and local edge at pixel $\boldsymbol{p}_M = \boldsymbol{p}_{15}$ are being checked.

4. Rectangle fitting

The edge map that is computed for every θ_n contains positive and negative edge points. Due to the handling of connectivity, the edge points of a line segment do not exhibit gaps between them and, consequently, a simple region growing scheme is sufficient to obtain the areas that form a rectangle. This avoids expensive search mechanisms for overcoming gaps, as the costly

Algorithm 3 mark connected edge points (assumes $H(\theta_n) = H$)

input: signed TV distance $\delta_{X_T X_B}$, directional data $\boldsymbol{\nabla}_{Q(\theta_n)}\boldsymbol{I}$, maximum distance threshold d, contextual C and local L_C edge thresholds
if $|\delta_{X_T X_B}| \geq C$ **then**
  if $\operatorname{sgn}(\delta_{X_T X_B})\, \boldsymbol{\nabla}_{Q(\theta_n)}(\boldsymbol{p}_M = \boldsymbol{A}(x + M - 1, y, \theta_n)) \geq L_C$ **then**
    % found valid local edge
    $g_C \leftarrow 0$
  else
    $g_C \leftarrow g_C + 1$
  end if
  if $g_C \leq d$ **then**
    % does not exceed d, mark $\boldsymbol{p}_M$ as a connected edge point
    $\boldsymbol{E}(\boldsymbol{p}_M) = \operatorname{sgn}(\delta_{X_T X_B})$
  else
    % unmark non-connected edge points
    for $m \in \{M - g_C + 1, \ldots, M\}$ **do**
      $\boldsymbol{E}(\boldsymbol{p}_m = \boldsymbol{A}(x + m - 1, y, \theta_n)) = 0$
    end for
    step = 'search for start'
  end if
else
  step = 'search for start'
end if

scheme used in our previous HT-based method, in [22]. This section works by:

- *1. Region growing and labeling* - A unique identification number is given to each set of connected edge points with the same sign, as illustrated in Fig. 8.

- *2. Fitting a rectangle to each area* - Rectangle fitting can occur in various ways (see [9] for a brief summary). In our method, we start by fitting a line to the upper and lower limits of each area, as illustrated in Fig. 8. Then, using the average angle of both fitted lines, $\theta_R = (\theta_{upper} + \theta_{lower})/2$, we obtain the start and end of each line segment.

- *3. Validate rectangles* - A validation step is needed to eliminate non-rectilinear structures in the image. Although various criteria can be used, we require only that the upper and lower fitted lines (see above) should have similar angles, $\theta_{upper} \sim \theta_{lower}$, and the average angle should lie inside the

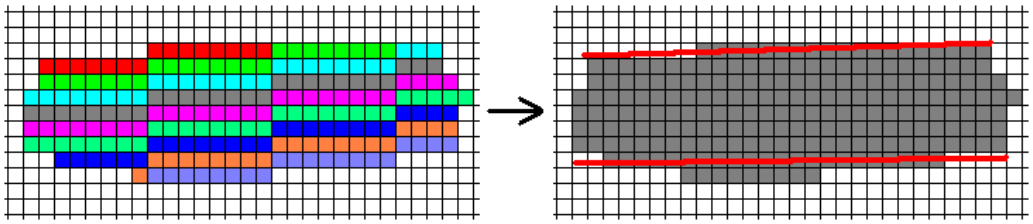

Figure 8: The multiple connected edges along lines (left), obtained in Section 3.2, are joined into a single area and lines are fit to the upper and lower limits (right).

permitted range, $\theta_R \in [\theta_n - 180°/2N, \theta_n + 180°/2N]$.
- *4. Store rectangles* - Valid rectangles are then stored as angle θ_R and the parameters corresponding to the four limiting lines, represented in Fig. 9.

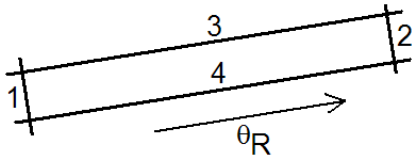

Figure 9: Rectangle limits.

5. Experiments

We single out demonstrative results of our method, which we contrast with the ones obtained with the standard HT [13], the state-of-the-art of local methods LSD [9] (the superiority of LSD when compared to several other local methods is thoroughly demonstrated in [9]), and our previous HT-based method, denoted STRAIGHT [22]. We describe experiments with synthetic images, which help characterize the general behavior of the new method. Then, we present results obtained with several real world images, which demonstrate its performance in practical application. Finally, we discuss the computational complexity of these methods.

5.1. Synthetic images

We start by illustrating the behavior of the algorithms when dealing with an image made up of intersecting line segments of multiple lengths and widths, shown on the top left of Fig. 10. By comparing the edges computed by the Canny edge detector [18] with the line segments that the HT

extracts from them, on the top middle and right of Fig. 10, respectively, we conclude that the HT succeeds in correctly extracting the lines from this image. This occurs because line segments are long, not in a large number, and the HT does not require connectivity, therefore being able to overcome the multiple line crossings. On the other hand, the results of the LSD method, shown in the bottom left image of Fig. 10, illustrate that local methods fail to overcome line crossings and splits them. This occurs because local methods require absolute connectivity, *i.e.*, that edge points are perfectly chained together. In the particular case of the LSD, the state-of-the-art of local methods, edge points must have approximately constant direction as well. Although the results of STRAIGHT show that it is able to extract thin line segments regardless of the intersections, it is unable to deal with thick ones, originating multiple erroneous detections. Our semi-global method succeeds in extracting line segments of all lengths and widths, with few errors. A pair of twin segments is extracted for each segment in the original image because both light-to-dark and dark-to-light transitions are detected.

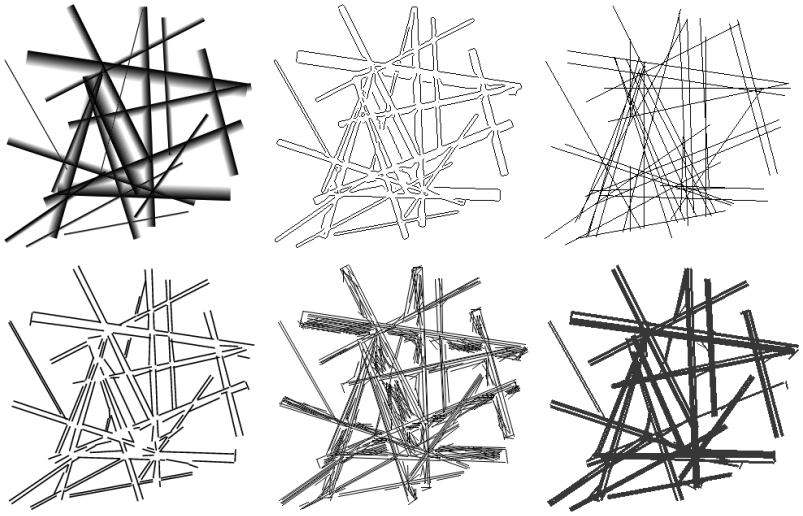

Figure 10: Image with prominent lines. Top left to right: original image, result of Canny edge detector [18], and standard HT [13]. Bottom left to right: LSD [9], STRAIGHT [22], and the proposed method.

We now illustrate the behavior of the algorithms in capturing transitions between differently textured regions. This simulates low signal-to-noise scenarios that occur when using very low thresholds in edge detection, for

increased sensibility, where real transitions should be extracted successfully, while avoiding false ones. We use the synthetic images in the left column of Fig. 11, which were generated by adding noise to a piecewise constant map (the top and bottom images were first used in [22] and are used here to enable a simple comparison between the methods). The top image represents a simpler scenario, where one of the areas involved in the transition is perfectly smooth. The central image represents the same scenario, except that the mean value of both regions is now equal, making the variance the single discriminating factor. The bottom image simulates two smooth objects in a low signal-to-noise image. Due to the effect of noise, the local edge detection in LSD can not produce edges with constant direction along real segments and the LSD fails to extract line segments, except for parts of the top image. The results of STRAIGHT and the proposed method, in the two rightmost columns of Fig. 11, show that both overcome noise and succeed in extracting the line segments for the top and bottom images (the few short segments correspond to accidental connected alignments in the random texture). By allowing samples with the same mean but different variances to be classified as edges, by using two-sample tests and, in particular, the TV distance, the proposed method is the only one that succeeds in obtaining most of the real segments of the figure in the middle row.

5.2. Real images

We start by showing a challenging image that was first used in [22] to demonstrate the ability of STRAIGHT in dealing with the dense packing of line segments of multiple lengths that cross each other. In the top right image of Fig. 12, we display the results of the HT [13], showing that extraction fails altogether for not being able to cope with the large number of edge points (this is explained in detail in [22]). On the middle left image of Fig. 12, we display the results of LSD [9], showing that a subset of the line segments are in fact detected. A closer look reveals that those are only the line segments that do not cross other structures and also that several longer segments are detected as fragmented ones. The results of STRAIGHT are shown in the middle right image of Fig. 12 and we see that it succeeds in extracting the vast majority of the line segments in the image (exceptions are those which exhibit low contrast). Equally good results are obtained by the method we propose, shown in the bottom images of Fig. 12, with the difference that STRAIGHT needed 52 seconds to process this image while our semi-global method needed only about 4 seconds on the same machine. The bottom

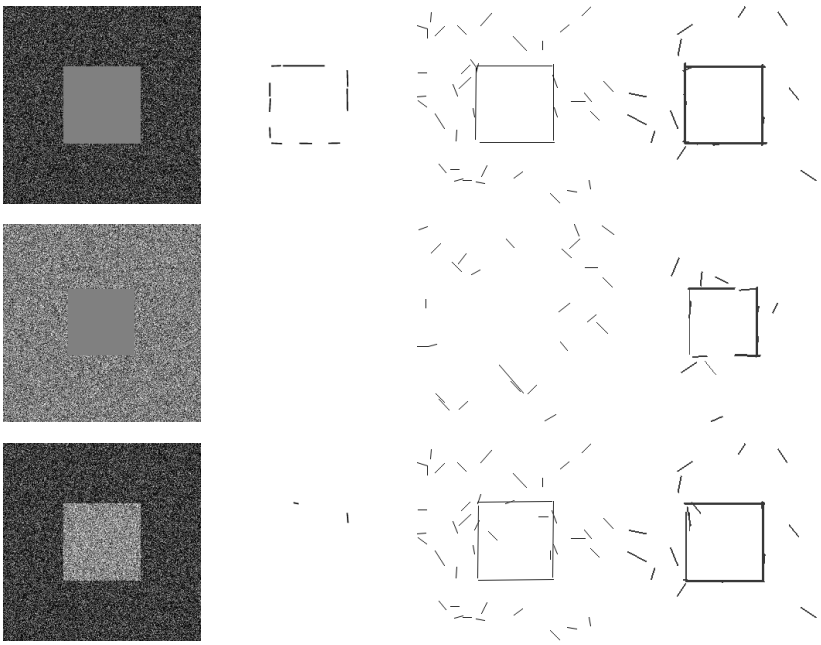

Figure 11: Textured images. From left to right: original image, result of LSD [9], STRAIGHT [22] and the proposed method.

right image displays only the line segments that have length greater than 50 pixels, illustrating that line segments are not fragmented.

Fig. 13 presents another illustrative case. It was obtained by processing an image containing a complex scene of line segments (many of which of low contrast) occluded by a net that is large and out-of-focus. The result of LSD [9] shows that most low contrast segments were not extracted and that others are fragmented in multiple pieces. The fragmentation of line segments is improved in STRAIGHT [22] but low contrast segments are equally not extracted and the thick lines of the net originate a multitude of erroneous line segments. On the other hand, our proposed method extracts most line segments, including low contrast and thick ones, with little fragmentation. The extraction of low contrast line segments is enabled by the better handling of variable high frequency variation by two-sample statistical tests. Low contrast and blurred line segments occur often in realistic scenarios and are tackled by our method.

Finally, Fig. 14 presents results of using our method with real images of various kinds. As desired, the vast majority of long line segments are ex-

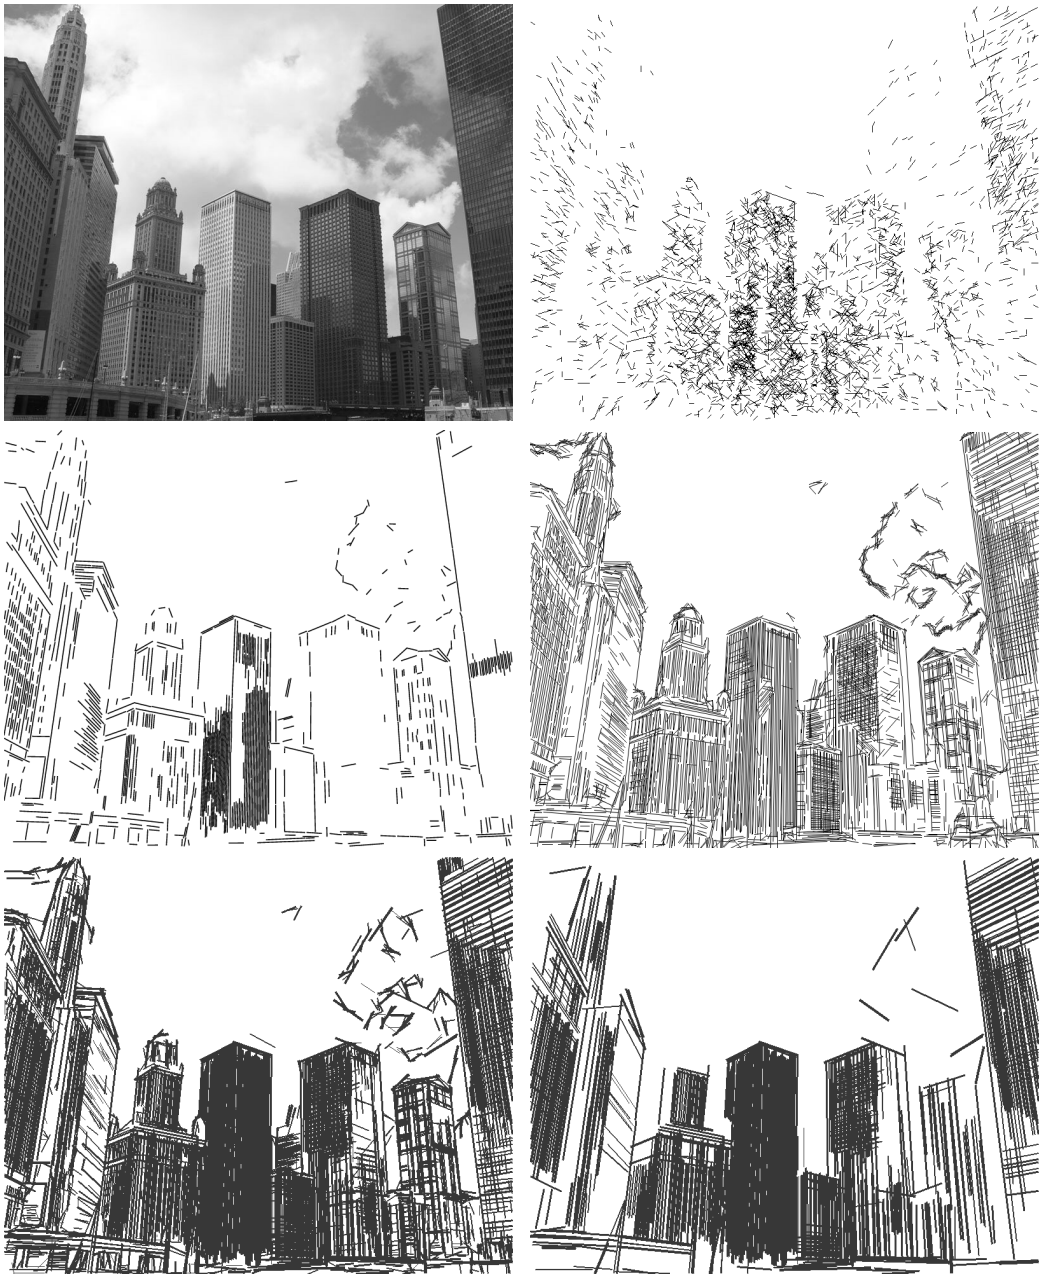

Figure 12: Top left: image. Top right: HT [13]. Middle left: LSD [9]. Middle right: STRAIGHT [22]. Bottom left: our method. Bottom right: our method (longer segments).

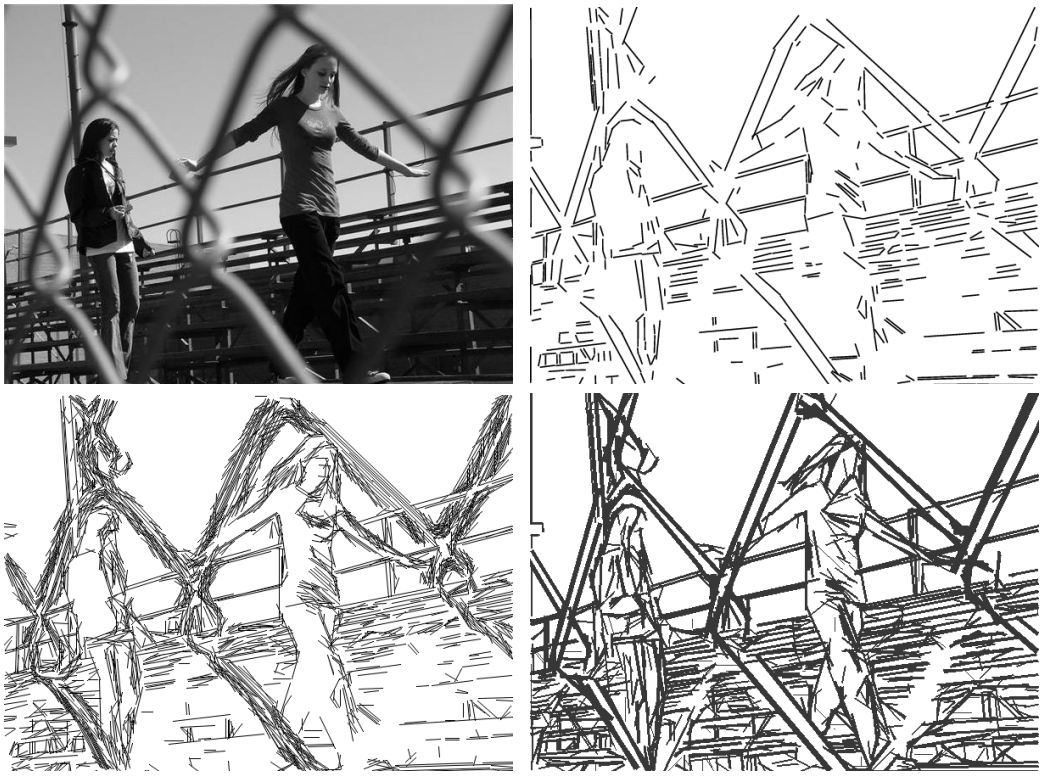

Figure 13: Top left: image. Top right: LSD [9]. Bottom left: STRAIGHT [22]. Bottom right: our method.

tracted without artificial fragmentation, despite the multiple segment crossings. Also note that, although some of these images have edges that form curves, our method succeeds in approximating these sections in a piecewise linear way, *i.e.*, by a sequence of rectilinear line segments.

5.3. Computational complexity

The most computationally intensive portion of our method is the calculation of contextual and local edges and their combination into continuous edge maps. By using a running average framework, the calculation of contextual edges depends only linearly on the pixel count and the number of directions, N. Such linear dependency also occurs in the calculation of local edges and in their combination with contextual ones. This is confirmed in

the left of Fig. 15, which shows the computation time[1] needed by the proposed method to extract line segments for multiple images, as a function of pixel count and the number of directions that are used, N. The trend line for $N = 32$, the number of directions that are used in all experiments, indicates that about 97K pixels are processed at each second (*e.g.*, the line segments in the 512×512 Lenna image are extracted in about 2.7 seconds).

The complexity of STRAIGHT [22] increases linearly with the number of edge points, as the dominating factor in the calculations is the updating of the Hough space of each local HT. In our experiments, the computation time of STRAIGHT was always considerably higher than the proposed method (*e.g.*, the line segments in the Lenna image are extracted in about 36.8 seconds), which is confirmed on the right of Fig. 15. This figure shows the computation time needed by both methods to extract line segments for multiple images, as a function of pixel count and for $N = 32$ directions. Although the performance of STRAIGHT varies non-linearly with the pixel count, the rough trendline that is fitted helps in illustrating the significant computational advantage of the method we propose of about one order of magnitude.

The standard HT requires filling an accumulator array, which depends linearly on the total number of edge points. In our experiments, the standard HT required approximately 1 to 10 seconds to process each image in this paper, using MATLAB© code.

Reference [9] states that the complexity of the LSD method depends linearly with the image pixel count, as illustrated by a plot therein showing the calculation time needed to extract line segments in various images. In the worst case scenario, *i.e.*, images made up of noise, LSD is able to process about 240K pixels per second. Although the main advantage of local methods such as the LSD is its low computational complexity, with the drawback of only dealing successfully with simple scenarios, the amount of pixels processed by the LSD is only about 2–3 times greater than our proposed method.

Although only STRAIGHT and our proposed method can deal with the complex images that arise in practice, the results above show that the computational complexity of our method is significantly below STRAIGHT. Fur-

[1] All experiments were performed on an Intel© 2.67 GHz machine, with all methods implemented in optimized C# code, except where noted.

thermore, the complexity of our method is comparable with local methods, *i.e.*, it is only about 2–3 times more complex than the LSD method, despite the ability to handle complex scenarios. This indicates that our method is efficient in extracting segments of all lengths and widths in complex scenarios.

6. Conclusion

We have presented a new semi-global method for line segment extraction. Our method combines contextual and local edges, with explicit handling of connectivity. Our experiments show that it outperforms current methods for line segment extraction in challenging situations, *e.g.*, when dealing with complex images containing several crossing segments of multiple widths, and that its computational efficiency is comparable with simple local methods. We use a contextual edge detection scheme based on two-sample statistical tests, which is a robust way to handle noise.

References

[1] J. Kosecka, W. Zhang, Video Compass, in: Proceedings of European Conference on Computer Vision, Springer-Verlag, Copenhagen, Denmark, 657–673, 2002.

[2] C. Schmid, A. Zisserman, Automatic Line Matching across Views, in: Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition, San Juan, Puerto Rico, 666–671, 1997.

[3] B. Micusk, H. Wildenauer, J. Kosecka, Detection and matching of rectilinear structures, in: Proceedings of IEEE International Conference on Computer Vision and Pattern Recognition, Anchorage, Alaska, USA, 1–7, 2008.

[4] D. Slater, G. Healey, 3D Shape Reconstruction by Using Vanishing Points, IEEE Transactions on Pattern Analysis and Machine Intelligence 18 (1996) 211–217.

[5] W. Krüger, Robust and efficient map-to-image registration with line segments, Machine Vision and Applications 13 (2001) 38–50.

[6] J. Liu, Y. Chen, X. Tang, Decomposition of Complex Line Drawings with Hidden Lines for 3D Planar-Faced Manifold Object Reconstruction, IEEE Transactions on Pattern Analysis and Machine Intelligence 33 (2011) 3–15.

[7] P. Fränti, E. Ageenko, H. Kälviäinen, S. Kukkonen, Compression of line drawing images using Hough transform for exploiting global dependencies, in: Proceedings of the Fourth Joint Conference on Information Sciences, vol. 4, Research Triangle Park, North Carolina, USA, 433–436, 1998.

[8] S. Du, B. J. van Wyk, C. Tu, X. Zhang, An improved Hough transform neighborhood map for straight line segments, IEEE Transactions on Image Processing 19 (2010) 573–585.

[9] R. von Gioi, J. Jakubowicz, J. Morel, G. Randall, LSD: A Fast Line Segment Detector with a False Detection Control, IEEE Transactions on Pattern Analysis and Machine Intelligence 32 (4) (2010) 722–732.

[10] A. Borkar, M. Hayes, M. Smith, Polar randomized Hough Transform for lane detection using loose contraints of parallel lines, in: Proceedings of IEEE International Conference on Acoustics, Speech, and Signal Processing, Prague, Czech Republic, 1037–1040, 2011.

[11] J. Ji, G. Chen, L. Sun, A novel Hough Transform method for line detection by enhancing accumulator array, Pattern Recognition Letters 32 (11) (2011) 1503–1510.

[12] P. Hough, Method and means for recognizing complex patterns, U.S. Patent 3.069.654, 1962.

[13] R. O. Duda, P. E. Hart, Use of the Hough Transformation to detect lines and curves in pictures., Communications of the ACM 15 (1) (1972) 11–15.

[14] O. Duprat, B. Keck, C. Ruwwe, U. Zölzer, Hough Transform with weighting Edge-maps, in: Fifth IASTED International Conference on Visualization, Imaging, & Image Processing, Benidorm, Spain, 2005.

[15] V. Kamat-Sadekar, S. Ganesan, Complete description of multiple line segments using the Hough Transform, Image and Vision Computing 16 (9-10) (1998) 597 – 613.

[16] M. Teutsch, T. Schamm, Fast Line and Object Segmentation in Noisy and Cluttered Environments using Relative Connectivity, in: Proceedings of the Conference on Image Processing, Computer Vision, and Pattern Recognition, Las Vegas, USA, 2011.

[17] A., A. Zeevi, The Hough Transform estimator, Annals of statistics 32 (2004) 1908.

[18] J. Canny, A Computational Approach To Edge Detection, IEEE Transactions on Pattern Analysis and Machine Intelligence 8 (1986) 679–698.

[19] M. Bertero, T. Pogcio, V. Torre, Ill-posed problems in early vision, in: Proceedings of the IEEE, 869–889, 1988.

[20] J. Kim, R. Krishnapuram, A robust Hough Transform based on validity, in: IEEE International Conference on Fuzzy Systems, vol. 2, Anchorage, USA, 1530 – 1535, 1998.

[21] N. Guil, J. Villalba, E. Zapata, A fast Hough Transform for segment detection., IEEE Transactions on Image Processing 4 (11) (1995) 1541–1548.

[22] R. Guerreiro, P. Aguiar, Connectivity-Enforcing Hough Transform for the Robust Extraction of Line Segments (to appear), IEEE Transactions on Image Processing .

[23] S. Yuen, T. Lam, N. Leung, Connective Hough transform, Image and Vision Computing 11 (5) (1993) 295–301.

[24] M. Yang, J. Lee, C. Lien, C. Huang, Hough Transform Modified by Line Connectivity and Line Thickness, IEEE Transactions on Pattern Analysis and Machine Intelligence 19 (8) (1997) 905–910.

[25] J. Illingworth, J. Kittler, A survey of efficient Hough Transform methods, in: Proceedings of Alvery Vision Club Conference, Cambridge, 319–326, 1987.

[26] J. Zhou, W. Bischof, A. Sanchez-Azofeifa, Extracting Lines in Noisy Image Using Directional Information, in: Proceedings of the 18th International Conference on Pattern Recognition, vol. 2, IEEE Computer Society, Washington, DC, USA, 215–218, 2006.

[27] H. Kälviäinen, P. Hirvonen, L. Xu, E. Oja, Probabilistic and non-probabilistic Hough Transforms: overview and comparisons, Image and Vision Computing 13 (1995) 239–252.

[28] H. Li, M. Lavin, R. Le Master, Fast Hough Transform: A hierarchical approach, Computer Vision, Graphics, and Image Processing 36 (1986) 139–161.

[29] J. Illingworth, J. Kittler, The adaptive Hough Transform, IEEE Transactions on Pattern Analysis and Machine Intelligence 9 (1987) 690–698.

[30] R. Stephens, Probabilistic approach to the Hough transform, Image and Vision Computing 9 (1) (1991) 66 – 71.

[31] R. Nevatia, K. Babu, Linear feature extraction and description, Computer Graphics and Image Processing 13 (3) (1980) 257 – 269.

[32] J. Burns, A. Hanson, E. Riseman, Extracting straight lines, IEEE Transactions on Pattern Analysis and Machine Intelligence 8 (1986) 425–455.

[33] D. Guru, A simple and robust line detection algorithm based on small eigenvalue analysis, Pattern Recognition Letters 25 (1) (2004) 1–13.

[34] V. Nguyen, S. Gächter, A. Martinelli, N. Tomatis, R. Siegwart, A comparison of line extraction algorithms using 2D range data for indoor mobile robotics, Autonomous Robots 23 (2007) 97–111.

[35] A. Etemadi, Robust Segmentation of Edge Data, in: Proceedings of IEEE International Conference on Image Processing and its Applications, Maastricht, The Netherlands, 311–314, 1992.

[36] O. Faugeras, R. Deriche, H. Mathieu, N. Ayache, G. Randall, Parallel image processing, chap. The depth and motion analysis machine, World Scientific Publishing Co., Inc., River Edge, NJ, USA, 143–175, 1992.

[37] K. Arras, R. Siegwart, Feature Extraction and Scene Interpretation for Map-Based Navigation and Map Building, in: Proceedings of SPIE, Mobile Robotics XII, vol. 3210, 42–53, 1997.

[38] J. Xiao, M. Shah, Two-Frame Wide Baseline Matching, in: Proceedings of IEEE International Conference on Computer Vision, Nice, France, 603–609, 2003.

[39] L. Wilson, N. Chen, R. Kelley, J. Birk, Image Feature Extraction Using Diameter Limited Gradient Direction Histograms, IEEE Transactions on Pattern Analysis and Machine Intelligence 2 (1) (1979) 228–235.

[40] A. Desolneux, L. Moisan, J. Morel, Gestalt theory and image analysis, a probabilistic approach, 2006.

[41] W. Freeman, E. Adelson, The Design and Use of Steerable Filters, IEEE Transactions on Pattern Analysis and Machine Intelligence 13 (1991) 891–906.

[42] A. Bovik, T. Huang, D. Munson, Nonparametric tests for edge detection in noise, Pattern Recognition 19 (3) (1986) 209–219.

[43] D. Sheskin, Handbook of Parametric and Nonparametric Statistical Procedures, Chapman & Hall/CRC, 4 edn., 2007.

[44] D. Lim, S. Jang, Comparison of two-sample tests for edge detection in noisy images, Journal of the Royal Statistical Society: Series D (The Statistician) 51 (1) (2002) 21–30.

[45] M. Fesharaki, G. Hellestrand, A new edge detection algorithm based on a statistical approach, in: Proceedings of International Symposium on Speech, Image Processing and Neural Networks, vol. 1, Hong Kong, China, 21 – 24, 1994.

[46] M. Thune, B. Olstad, N. Thune, Edge detection in noisy data using finite mixture distribution analysis, Pattern Recognition 30 (5) (1997) 685 – 699.

[47] I. Williams, D. Svoboda, N. Bowring, E. Guest, Improved statistical edge detection through neural networks, in: 10th Conference on Medical Image Understanding and Analysis, Manchester, UK, 56–60, 2006.

[48] A. DasGupta, Asymptotic theory of statistics and probability, Springer, 2010.

Figure 14: Results of our method for several kinds of real images.

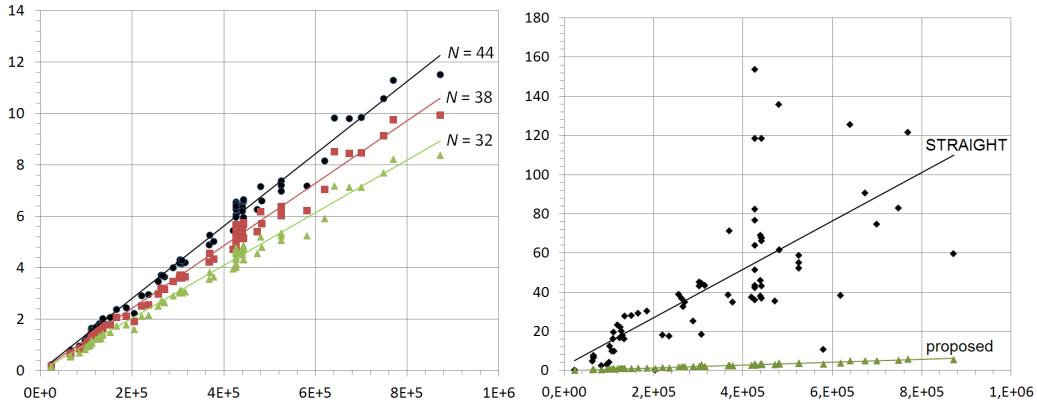

Figure 15: Extraction times (in seconds) as a function of pixel count. Left: proposed method, for N number of directions. About 71K pixels per second are processed using $N = 44$ directions, 82K pixels/s using $N = 38$, and 97K pixels/s using $N = 32$. Right: proposed method versus STRAIGHT, with $N = 32$. STRAIGHT processes roughly 8K pixels/s.